EXHIBITS A1-A6

(Part 7 of 13)

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
| snmp-server group — Configures a new SNMP group or a table that maps SNMP users to SNMP views.<br><br>snmp-server trap authentication vrf — Controls VRF-specific SNMP authentication failure notifications.<br><br>snmp-server user — Configures a new user to an SNMP group.<br><br>Cisco IOS SNMP Support Command Reference (2013), at 130. | **Configuring the Group**<br>An SNMP group is a table that maps SNMP users to SNMP views. The snmp-server group command configures a new SNMP group.<br><br>**Example**<br>• This command configures *normal_one* as an SNMPv3 group (authentication and encryption) that provides access to the *all-items* read view.<br>```
switch(config)#snmp-server group normal_one v3 priv read all-items
switch(config)#
```<br><br>**Configuring the User**<br>An SNMP user is a member of an SNMP group. The snmp-server user command adds a new user to an SNMP group and configures that user's parameters. To configure a remote user, specify the IP address or port number of the device where the user's remote SNMP agent resides.<br><br>Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1966.<br><br>*See also* Arista User Manual v. 4.13.6F (4/14/2014), at 1894; Arista User Manual v. 4.12.3 (7/17/13), at 1656; Arista User Manual, v. 4.11.1 (1/11/13), at 1343-44; Arista User Manual v. 4.10.3 (10/22/12), at 1109-10; Arista User Manual v. 4.9.3.2 (5/3/12), at 865; Arista User Manual v. 4.8.2 (11/18/11), at 677; Arista User Manual v. 4.7.3 (7/18/11), at 533. | Dkt. 419-10 at PDF p. 212 |

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record

snmp trap link-status
 	
To enable Simple Network Management Protocol (SNMP) link trap generation, use the **snmp trap link-status** command in either interface configuration mode or service instance configuration mode. To disable SNMP link trap generation, use the **no** form of this command.
 	
snmp trap link-status [permit duplicates]
no snmp trap link-status [permit duplicates]
 	
Cisco IOS SNMP Support Command Reference (2013), at 130.
	snmp trap link-status
	
	The snmp trap link-status command enables Simple Network Management Protocol (SNMP) link-status trap generation on the configuration mode interface. The generation of link-status traps is enabled by default. If SNMP link-trap generation was previously disabled, this command removes the corresponding no snmp link-status statement from the configuration to re-enable link-trap generation.
	
	The no snmp trap link-status command disables SNMP link trap generation on the configuration mode interface.
	
	The snmp trap link-status and default snmp trap link-status commands restore the default behavior by removing the no snmp trap link-status command from *running-config*. Only the no form of this command is visible in *running-config*.
	
	Platform all
	Command Mode Interface-Ethernet Configuration
	           Interface-Loopback Configuration
	           Interface-Management Configuration
	           Interface-Port-channel Configuration
	           Interface-VLAN Configuration
	           Interface-VXLAN Configuration
	
	Command Syntax
	    `snmp trap link-status`
	    `no snmp trap link-status`
	    `default snmp trap link-status`
	
	Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1966.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 1692; Arista User Manual, v. 4.11.1 (1/11/13), at 1377; Arista User Manual v. 4.10.3 (10/22/12), at 1144; Arista User Manual v. 4.9.3.2 (5/3/12), at 898; Arista User Manual v. 4.8.2 (11/18/11), at 705; Arista User Manual v. 4.7.3 (7/18/11), at 561.
		Dkt. 419-10 at PDF p. 213

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
snmp-server host　Specifies the targeted recipient of an SNMP notification operation. Cisco IOS SNMP Support Command Reference (2013), at 191.	**Configuring the Host** The snmp-server host command specifies the recipient of a SNMP notification. An SNMP host is the recipient of an SNMP trap operation. The snmp-server host command sets the community string if it was not previously configured. Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1967. *See also* Arista User Manual v. 4.13.6F (4/14/2014), at 1895; Arista User Manual v. 4.12.3 (7/17/13), at 1656; Arista User Manual, v. 4.11.1 (1/11/13), at 1344; Arista User Manual v. 4.10.3 (10/22/12), at 1110; Arista User Manual v. 4.9.3.2 (5/3/12), at 866; Arista User Manual v. 4.8.2 (11/18/11), at 677; Arista User Manual v. 4.7.3 (7/18/11), at 533.	Dkt. 419-10 at PDF p. 214
Usage Guidelines　SNMP notifications can be sent as traps or inform requests. This command enables both traps and inform requests. Cisco IOS SNMP Support Command Reference (2013), at 216.	The snmp-server enable traps command enables the transmission of Simple Network Management Protocol (SNMP) notifications as traps or inform requests. This command enables both traps and inform requests for the specified notification types. The snmp-server host command specifies the notification type (traps or informs). Sending notifications requires at least one snmp-server host command. Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1990. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1680; Arista User Manual, v. 4.11.1 (1/11/13), at 1365; Arista User Manual v. 4.10.3 (10/22/12), at 1132; Arista User Manual v. 4.9.3.2 (5/3/12), at 888; Arista User Manual v. 4.8.2 at 696; Arista User Manual v. 4.7.3 (7/18/11), at 552.	Dkt. 419-10 at PDF p. 214

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
snmp-server engineID local To specify the Simple Network Management Protocol (SNMP) engine ID on the local device, use the **snmp-server engineID local** command in global configuration mode. To remove the configured engine ID, use the **no** form of this command. **snmp-server engineID local** *engineid-string* **no snmp-server engineID local** *engineid-string* **Syntax Description** | *engineid-string*	String of a maximum of 24 characters that identifies the engine ID. | **Command Default** An SNMP engine ID is generated automatically but is not displayed or stored in the running configuration. You can display the default or configured engine ID by using the **show snmp engineID** command. **Command Modes** Global configuration (config) **Command History** | Release	Modification
---	---
12.0(3)T	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware. | **Usage Guidelines** The SNMP engine ID is a unique string used to identify the device for administrative purposes. You do not need to specify an engine ID for the device; a default string is generated using Cisco's enterprise number (1.3.6.1.4.1.9) and the MAC address of the first interface on the device. For further details on the SNMP engine ID, see RFC 2571. If you specify your own ID, note that the entire 24-character engine ID is not needed if it contains trailing zeros. Specify only the portion of the engine ID up until the point where only zeros remain in the value. For example, to configure an engine ID of 123400000000000000000000, you can specify **snmp-server engineID local 1234**. The value for the engine ID is displayed in hexadecimal value pairs. If the length of the input is an odd number, the last digit will be prepended with a zero ("0"). For example, if the engine ID is 12345, the ID is treated as 12:34:05 internally. Hence, the engine ID is displayed as 123405 in the **show running configuration** command output. Changing the value of the SNMP engine ID has significant effects. A user's password (entered on the command line) is converted to a message digest5 algorithm (MD5) or Secure Hash Algorithm (SHA) security digest. This digest is based on both the password and the local engine ID. The command line password is then destroyed, as required by RFC 2274. Because of this deletion, if the local value of the engineID changes, the security digests of SNMPv3 users will become invalid, and the users will have to be reconfigured. Similar restrictions require the reconfiguration of community strings when the engine ID changes. A remote engine ID is required when an SNMPv3 inform is configured. The remote engine ID is used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host. **Examples** The following example specifies the local SNMP engine ID: `Router(config)# snmp-server engineID local` Cisco IOS SNMP Support Command Reference (2013), at 339-340.	**snmp-server engineID local** The **snmp-server engineID local** command configures the name for the local Simple Network Management Protocol (SNMP) engine. The default SNMP engineID is generated by the switch and is used when an engineID is not configured with this command. The **show snmp engineID** command displays the default or configured engine ID. SNMPv3 authenticates users through security digests (MD5 or SHA) that are based on user passwords and the local engine ID. Passwords entered on the CLI are similarly converted, then compared to the user's security digest to authenticate the user. **Important** Changing the local engineID value invalidates SNMPv3 security digests, requiring the reconfiguration of all user passwords. The **no snmp-server engineID** and **default snmp-server engineID** commands restore the default engineID by removing the **snmp-server engineID** command from the configuration. | Platform	all

Command Mode	Global Configuration |
	
	Command Syntax
	
	`snmp-server engineID local engine_hex`
	`no snmp-server engineID local`
	`default snmp-server engineID`
	
	Parameters
	• *engine_hex* the switch's name for the local SNMP engine (hex string).
	  The string must consist of at least ten characters with a maximum of 64 characters.
	
	Example
	• This command configures DC945798CAB4 as the name of the local SNMP engine.
	
	`switch(config)#snmp-server engineID local DC945798CAB4`
	`switch(config)#`
	
	snmp-server engineID remote
	
	The **snmp-server engineID remote** command configures the name of a Simple Network Management Protocol (SNMP) engine located on a remote device. The switch generates a default engineID; use the **show snmp engineID** command to view the configured or default engineID.
	
	A remote engine ID is required when configuring an SNMPv3 inform to compute the security digest for authenticating and encrypting packets sent to users on the remote host. SNMPv3 authenticates users through security digests (MD5 or SHA) that are based on user passwords and the engine ID. Passwords entered on the CLI are similarly converted, then compared to the user's security digest to authenticate the user.
	
	Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1991-92.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 1681-82; Arista User Manual, v. 4.11.1 (1/11/13), at 1366-67; Arista User Manual v. 4.10.3 (10/22/12), at 1133-34; Arista User Manual v. 4.9.3.2 (5/3/12), at 889-890; Arista User Manual v. 4.8.2 (11/18/11), at 697-98; Arista User Manual v.
		Dkt. 419-10 at PDF pp. 215-216

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
	4.7.3 (7/18/11), at 553-54.	
show snmp engineID / Displays the identification of the local SNMP engine and all remote engines that have been configured on the router. Cisco IOS SNMP Support Command Reference (2013), at 340/	**show snmp engineID** The show snmp engineID command displays the identification of the local Simple Network Management Protocol (SNMP) engine and of all remote engines that are configured on the switch. Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1978. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1668; Arista User Manual, v. 4.11.1 (1/11/13), at 1355; Arista User Manual v. 4.10.3 (10/22/12), at 1122; Arista User Manual v. 4.9.3.2 (5/3/12), at 878; Arista User Manual v. 4.8.2 (11/18/11), at 686; Arista User Manual v. 4.7.3 (7/18/11), at 542.	Dkt. 419-10 at PDF p. 216

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
snmp-server group To configure a new Simple Network Management Protocol (SNMP) group, use the **snmp-server group** command in global configuration mode. To remove a specified SNMP group, use the **no** form of this command. **snmp-server group** *group-name* {**v1**| **v2c**| **v3** {**auth**| **noauth**| **priv**}} [**context** *context-name*] [**read** *read-view*] [**write** *write-view*] [**notify** *notify-view*] [**access** [**ipv6** *named-access-list*] [*acl-number*| *acl-name*]] **no snmp-server group** *group-name* {**v1**| **v2c**| **v3** {**auth**| **noauth**| **priv**}} [**context** *context-name*] Syntax Description 	*group-name*	Name of the group.
v1	Specifies that the group is using the SNMPv1 security model. SNMPv1 is the least secure of the possible SNMP security models.
v2c	Specifies that the group is using the SNMPv2c security model. The SNMPv2c security model allows informs to be transmitted and supports 64-character strings.
v3	Specifies that the group is using the SNMPv3 security model. SMNPv3 is the most secure of the supported security models. It allows you to explicitly configure authentication characteristics.
auth	Specifies authentication of a packet without encrypting it.
noauth	Specifies no authentication of a packet.
priv	Specifies authentication of a packet with encryption.
context	(Optional) Specifies the SNMP context to associate with this SNMP group and its views.
context-name	(Optional) Context name.
read	(Optional) Specifies a read view for the SNMP group. This view enables you to view only the contents of the agent.		**snmp-server group** The snmp-server group command configures a new Simple Network Management Protocol (SNMP) group or modifies an existing group. An SNMP group is a data structure that user statements reference to map SNMP users to SNMP contexts and views, providing a common access policy to the specified users. An SNMP context is a collection of management information items accessible by an SNMP entity. Each item of may exist in multiple contexts. Each SNMP entity can access multiple contexts. A context is identified by the EngineID of the hosting device and a context name. The no snmp-server group and default snmp-server group commands delete the specified group by removing the corresponding snmp-server group command from the configuration. Platform all Command Mode Global Configuration Command Syntax ```\nsnmp-server group group_name VERSION [CNTX] [READ] [WRITE] [NOTIFY]\nno snmp-server group group_name VERSION\ndefault snmp-server group group_name VERSION\n``` Parameters • *group_name* the name of the group. • *VERSION* the security model used by the group. — **v1** SNMPv1. Uses a community string match for authentication. — **v2c** SNMPv2c. Uses a community string match for authentication. — **v3 no auth** SNMPv3. Uses a username match for authentication. — **v3 auth** SNMPv3. HMAC-MD5 or HMAC-SHA authentication. — **v3 priv** SNMPv3. HMAC-MD5 or HMAC-SHA authentication. AES or DES encryption. • *CNTX* associates the SNMP group to an SNMP context. — \<no parameter> command does not associate group with an SNMP context. — **context** *context_name* associates group with context specified by *context_name*. • *READ* specifies read view for SNMP group. — \<no parameter> command does not specify read view. — **read** *read_name* read view specified by *read_name* (string – maximum 64 characters). • *WRITE* specifies write view for SNMP group. — \<no parameter> command does not specify write view. — **write** *write_name* write view specified by *write_name* (string – maximum 64 characters). • *NOTIFY* specifies notify view for SNMP group. — \<no parameter> command does not specify notify view. — **notify** *notify_name* notify view specified by *notify_name* (string – maximum 64 characters).	Dkt. 419-10 at PDF p. 217

Cisco's Documentation		Arista's Documentation	Supporting Evidence In The Record
read-view	(Optional) String of a maximum of 64 characters that is the name of the view. The default is that the read-view is assumed to be every object belonging to the Internet object identifier (OID) space (1.3.6.1), unless the **read** option is used to override this state.	Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1994. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1684; Arista User Manual, v. 4.11.1 (1/11/13), at 1369; Arista User Manual v. 4.10.3 (10/22/12), at 1136; Arista User Manual v. 4.9.3.2 (5/3/12), at 892; Arista User Manual v. 4.8.2 (11/18/11), at 699; Arista User Manual v. 4.7.3 (7/18/11), at 555.	Dkt. 419-10 at PDF p. 218
write	(Optional) Specifies a write view for the SNMP group. This view enables you to enter data and configure the contents of the agent.		
write-view	(Optional) String of a maximum of 64 characters that is the name of the view. The default is that nothing is defined for the write view (that is, the null OID). You must configure write access.		
notify	(Optional) Specifies a notify view for the SNMP group. This view enables you to specify a notify, inform, or trap.		
notify-view	(Optional) String of a maximum of 64 characters that is the name of the view. By default, nothing is defined for the notify view (that is, the null OID) until the **snmp-server host** command is configured. If a view is specified in the **snmp-server group** command, any notifications in that view that are generated will be sent to all users associated with the group (provided a SNMP server host configuration exists for the user). Cisco recommends that you let the software autogenerate the notify view. See the "Configuring Notify Views" section in this document.		
access	(Optional) Specifies a standard access control list (ACL) to associate with the group.		
ipv6	(Optional) Specifies an IPv6 named access list. If both IPv6 and IPv4 access lists are indicated, the IPv6 named access list must appear first in the list.		
named-access-list	(Optional) Name of the IPv6 access list.		
acl-number	(Optional) The *acl-number* argument is an integer from 1 to 99 that identifies a previously configured standard access list.		

Cisco IOS SNMP Support Command Reference (2013), at 343-44.

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
■ **snmp-server host** 	Release	Modification
Cisco IOS XE Release 2.1	This command was integrated into Cisco IOS XE Release 2.1.

15.2(1)S	This command was modified. The **p2mp-traffic-eng** notification-type keyword was added. |
	
	Usage Guidelines If you enter this command with no optional keywords, the default is to send all notification-type traps to the host. No informs will be sent to the host.
	
	The **no snmp-server host** command with no keywords disables traps, but not informs, to the host. To disable informs, use the **no snmp-server host informs** command.
	
	Note If a community string is not defined using the snmp-server community command prior to using this command, the default form of the **snmp-server community** command will automatically be inserted into the configuration. The password (community string) used for this automatic configuration of the **snmp-server community** command will be the same as that specified in the **snmp-server host** command. This automatic command insertion and use of passwords is the default behavior for Cisco IOS Release 12.0(3) and later releases. However, in Cisco IOS Release 12.2(33)SRE and later releases, you must manually configure the **snmp-server community** command. That is, the **snmp-server community** command will not be seen in the configuration.
	
	Cisco IOS SNMP Support Command Reference (2013), at 354.
		snmp-server host
		
		The snmp-server host command specifies the recipient of Simple Network Management Protocol (SNMP) notifications. Recipients are denoted by host location and community string. The command also specifies the type of SNMP notifications that are sent: a *trap* is an unsolicited notification; an *inform* is a trap that includes a request for a confirmation that the message is received.
		
		The configuration can contain multiple statements to the same host location with different community strings. For instance, a configuration can simultaneously contain all of the following:
		
		• snmp-server host host-1 version 2c comm-1
		• snmp-server host host-1 informs version 2c comm-2
		• snmp-server host host-1 version 2c comm-3 udp-port 666
		• snmp-server host host-1 version 3 auth comm-3
		
		The no snmp-server host and default snmp-server host commands remove the specified host by deleting the corresponding snmp-server host statement from the configuration. When removing a statement, the host (address and port) and community string must be specified.
		
		Platform all
		Command Mode Global Configuration
		
		Command Syntax
		`snmp-server host host_id [VRF_INST] [MESSAGE] [VERSION] comm_str [PORT]`
		`no snmp-server host host_id [VRF_INST] [MESSAGE] [VERSION] comm_str [PORT]`
		`default snmp-server host host_id [VRF_INST] [MESSAGE] [VERSION] comm_str [PORT]`
		
		Parameters
		• *host_id* hostname or IP address of the targeted recipient.
		• *VRF_INST* specifies the VRF instance being modified.
		— <no parameter> changes are made to the default VRF.
		— vrf *vrf_name* changes are made to the specified user-defined VRF.
		• *MESSAGE* message type that is sent to the host.
		— <no parameter> sends SNMP traps to host (default).
		— informs sends SNMP informs to host.
		— traps sends SNMP traps to host.
		• *VERSION* SNMP version. Options include:
		— <no parameter> SNMPv2c (default).
		— version 1 SNMPv1; option not available with informs.
		— version 2c SNMPv2c.
		— version 3 noauth SNMPv3; enables user-name match authentication.
		— version 3 auth SNMPv3; enables MD5 and SHA packet authentication.
		— version 3 priv SNMPv3. HMAC-MD5 or HMAC-SHA authentication. AES or DES encryption.
		• *comm_str* community string (used as password) sent with the notification operation.
		Although this string can be set with the snmp-server host command, the preferred method is defining it with the snmp-server community command prior to using this command.
		
		Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1995.
		
		See also Arista User Manual v. 4.12.3 (7/17/13), at 1685; Arista User
			Dkt. 419-10 at PDF pp. 219-220

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
	Manual, v. 4.11.1 (1/11/13), at 1370; Arista User Manual v. 4.10.3 (10/22/12), at 1137; Arista User Manual v. 4.9.3.2 (5/3/12), at 893; Arista User Manual v. 4.8.2 (11/18/11), at 700; Arista User Manual v. 4.7.3 (7/18/11), at 556.	
SNMP notifications can be sent as traps or inform requests. Traps are unreliable because the receiver does not send acknowledgments when it receives traps. The sender cannot determine if the traps were received. However, an SNMP entity that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the sender never receives the response, the inform request can be sent again. Thus, informs are more likely to reach their intended destination than traps. Compared to traps, informs consume more resources in the agent and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once; an inform may be tried several times. The retries increase traffic and contribute to a higher overhead on the network. Cisco IOS SNMP Support Command Reference (2013), at 354.	SNMP notifications are messages, sent by the agent, to inform managers of an event or a network condition. A *trap* is an unsolicited notification. An *inform* (or inform request) is a trap that includes a request for a confirmation that the message is received. Events that a notification can indicate include improper user authentication, restart, and connection losses. Traps are less reliable than informs because the receiver does not send any acknowledgment. However, traps are often preferred because informs consume more switch and network resources. A trap is sent only once and is discarded as soon as it is sent. An inform request remains in memory until a response is received or the request times out. An inform may be retried several times, increasing traffic and contributing to higher network overhead. Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1963. *See also* Arista User Manual v. 4.13.6F (4/14/2014), at 1891; Arista User Manual v. 4.12.3 (7/17/13), at 1653; Arista User Manual, v. 4.11.1 (1/11/13), at 1341; Arista User Manual v. 4.10.3 (10/22/12), at 1107; Arista User Manual v. 4.9.3.2 (5/3/12), at 863; Arista User Manual v. 4.8.2 at 675; Arista User Manual v. 4.7.3 (7/18/11), at 531.	Dkt. 419-10 at PDF p. 220

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record

snmp-server source-interface
 	
 To specify the interface from which a Simple Network Management Protocol (SNMP) trap originates the informs or traps, use the **snmp-server source-interface** command in global configuration mode. To remove the source designation, use the **no** form of this command.
 	
 snmp-server source-interface {traps| informs} *interface*
 no snmp-server source-interface {traps| informs} [*interface*]
 	
 Cisco IOS SNMP Support Command Reference (2013), at 376.
	snmp-server source-interface
	
	 The snmp-server source-interface command specifies the interface from which a Simple Network Management Protocol (SNMP) trap originates the informs or traps.
	
	 The no snmp-server source-interface and default snmp-server source-interface commands remove the inform or trap source assignment by removing the snmp-server source-interface command from running-config.
	
	 Platform all
	 Command Mode Global Configuration
	
	 Command Syntax
	 `snmp-server source-interface INTERFACE`
	 `no snmp-server source-interface`
	 `default snmp-server source-interface`
	
	 Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1967.
	
	 See also Arista User Manual v. 4.12.3 (7/17/13), at 1688; Arista User Manual, v. 4.11.1 (1/11/13), at 1373; Arista User Manual v. 4.10.3 (10/22/12), at 1140; Arista User Manual v. 4.9.3.2 (5/3/12), at 895; Arista User Manual v. 4.8.2 (11/18/11), at 702; Arista User Manual v. 4.7.3 (7/18/11), at 558.
		Dkt. 419-10 at PDF p. 221

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
snmp-server user To configure a new user to a Simple Network Management Protocol (SNMP) group, use the **snmp-server user** command in global configuration mode. To remove a user from an SNMP group, use the **no** form of this command. **snmp-server user** *username group-name* [**remote** *host* [**udp-port** *port*] [**vrf** *vrf-name*]] {**v1**| **v2c**| **v3** [**encrypted**] [**auth** {**md5**| **sha**} *auth-password*]} [**access** [**ipv6** *nacl*] [**priv** {**des**| **3des**| **aes** {**128**| **192**| **256**}} *privpassword*] {*acl-number*| *acl-name*}] **no snmp-server user** *username group-name* [**remote** *host* [**udp-port** *port*] [**vrf** *vrf-name*]] {**v1**| **v2c**| **v3** [**encrypted**] [**auth** {**md5**| **sha**} *auth-password*]} [**access** [**ipv6** *nacl*] [**priv** {**des**| **3des**| **aes** {**128**| **192**| **256**}} *privpassword*] {*acl-number*| *acl-name*}] **Syntax Description** *username* — Name of the user on the host that connects to the agent. *group-name* — Name of the group to which the user belongs. **remote** — (Optional) Specifies a remote SNMP entity to which the user belongs, and the hostname or IPv6 address or IPv4 IP address of that entity. If both an IPv6 address and IPv4 IP address are being specified, the IPv6 host must be listed first. *host* — (Optional) Name or IP address of the remote SNMP host. **udp-port** — (Optional) Specifies the User Datagram Protocol (UDP) port number of the remote host. *port* — (Optional) Integer value that identifies the UDP port. The default is 162. **vrf** — (Optional) Specifies an instance of a routing table. *vrf-name* — (Optional) Name of the Virtual Private Network (VPN) routing and forwarding (VRF) table to use for storing data. **v1** — Specifies that SNMPv1 should be used. **v2c** — Specifies that SNMPv2c should be used. **v3** — Specifies that the SNMPv3 security model should be used. Allows the use of the **encrypted** keyword or **auth** keyword or both. Cisco IOS SNMP Support Command Reference (2013), at 394.	**snmp-server user** The snmp-server user command adds a user to a Simple Network Management Protocol (SNMP) group or modifies an existing user's parameters. To configure a remote user, specify the IP address or port number of the device where the user's remote SNMP agent resides. A remote agent's engine ID must be configured before remote users for that agent are configured. A user's authentication and privacy digests are derived from the engine ID and the user's password. The configuration command fails if the remote engine ID is not configured first. The no snmp-server user and default snmp-server user commands remove the user from an SNMP group by deleting the user command from *running-config*. Platform all Command Mode Global Configuration **Command Syntax** `snmp-server user` *user_name group_name* [*AGENT*] *VERSION* [*ENGINE*] [*SECURITY*] `no snmp-server user` *user_name group_name* [*AGENT*] *VERSION* `default snmp-server user` *user_name group_name* [*AGENT*] *VERSION* **Parameters** • *user_name* — name of the user on the host that connects to the agent. • *group_name* — name of the group to which the user is associated. • *AGENT* — location of the host connecting to the SNMP agent. Configuration options include: — <no parameter> local SNMP agent. — **remote** *addr* [**udp-port** *p_num*] remote SNMP agent location (IP address, udp port). *addr* denotes the IP address; *p_num* denotes the udp port socket. (default port is 162). • *VERSION* — SNMP version; options include: — **v1** SNMPv1. — **v2c** SNMPv2c. — **v3** SNMPv3; enables user-name match authentication. • *ENGINE* — engine ID used to localize passwords. Available only if *VERSION* is v3. — <no parameter> Passwords localized by SNMP copy specified by *agent*. — localized *engineID* octet string of engineID. • *SECURITY* — Specifies authentication and encryption levels. Available only if *VERSION* is v3. Encryption is available only when authentication is configured. — <no parameter> no authentication or encryption. — **auth** *a_meth a_pass* [**priv** *e_meth e_pass*] authentication and encryption parameters. *a-meth* authentication method: options are md5 (HMAC-MD5-96) and sha (HMAC-SHA-96). *a-pass* authentication string for users receiving packets. *e-meth* encryption method: tions are aes (AES-128) and des (CBC-DES). *e-pass* encryption string for the users sending packets. Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1999. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1689; Arista User	Dkt. 419-10 at PDF pp. 222-223

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
	Manual, v. 4.11.1 (1/11/13), at 1374; Arista User Manual v. 4.10.3 (10/22/12), at 1141; Arista User Manual v. 4.9.3.2 (5/3/12), at 896; Arista User Manual v. 4.8.2 (11/18/11), at 703; Arista User Manual v. 4.7.3 (7/18/11), at 559.	
Usage Guidelines To configure a remote user, specify the IP address or port number for the remote SNMP agent of the device where the user resides. Also, before you configure remote users for a particular agent, configure the SNMP engine ID, using the **snmp-server engineID** command with the **remote** keyword. The remote agent's SNMP engine ID is needed when computing the authentication and privacy digests from the password. If the remote engine ID is not configured first, the configuration command will fail. For the *privpassword* and *auth-password* arguments, the minimum length is one character; the recommended length is at least eight characters, and should include both letters and numbers. Cisco IOS SNMP Support Command Reference (2013), at 396.	To configure a remote user, specify the IP address or port number of the device where the user's remote SNMP agent resides. A remote agent's engine ID must be configured before remote users for that agent are configured. A user's authentication and privacy digests are derived from the engine ID and the user's password. The configuration command fails if the remote engine ID is not configured first. Arista User Manual v. 4.14.3F (Rev. 2) 10/2/2014), at 1999. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1689; Arista User Manual, v. 4.11.1 (1/11/13), at 1374; Arista User Manual v. 4.10.3 (10/22/12), at 1141; Arista User Manual v. 4.9.3.2 (5/3/12), at 896; Arista User Manual v. 4.8.2 (11/18/11), at 703; Arista User Manual v. 4.7.3 (7/18/11), at 559.	Dkt. 419-10 at PDF p. 223

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
timers basic (ISO CLNS) To configure ISO IGRP timers, use the **timers basic** command in router configuration mode. To restore the default values, use the **no** form of this command. **timers basic** *update-interval holddown-interval invalid-interval* **no timers basic** *update-interval holddown-interval invalid-interval* **Syntax Description** *update-interval* — Time, in seconds, between the sending of routing updates. *holddown-interval* — Time, in seconds, a system or area router is kept in holddown state, during which routing information regarding better paths is suppressed. (A router enters into a holddown state when an update packet is received that indicates the route is unreachable. The route is marked inaccessible and advertised as unreachable. However, the route is still used for forwarding packets.) When the holddown interval expires, routes advertised by other sources are accepted and the route is no longer inaccessible. *invalid-interval* — Time, in seconds, that a route remains in the routing table after it has been determined that it is not reachable. After that length of time, the route is removed from the routing table. Cisco IOS Interface and Hardware Component Command Reference (2011), at ISO-178.	**timers basic (RIP)** The timers basic command configures the update interval, the expiration time, and the deletion time for routes received and sent through RIP. The command requires value declaration of all values. • The update time is the interval between unsolicited route responses. The default is 30 seconds. • The expiration time is initialized when a route is established and any time an update is received for the route. If the specified period elapses from the last time the route update was received, then the route is marked as inaccessible and advertised as unreachable. However, the route forwards packets until the deletion time expires. The default value is 180 seconds. • The deletion time is initialized when the expiration time has elapsed. On initialization of the deletion time, the route is no longer valid; however, it is retained in the routing table for a short time so that neighbors can be notified that the route has been dropped. Upon expiration of the deletion time, the route is removed from the routing table. The default is 120 seconds. The no timers basic and default timers basic commands return the timer values to their default values by removing the timers-basic command from *running-config*. Platform all Command Mode Router-RIP Configuration **Command Syntax** **timers basic** *update* time expire_time deletion_time **no timers basic** default timers basic Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1671. *See also* Arista User Manual v. 4.13.6F (4/14/2014), at 1621; Arista User Manual v. 4.12.3 (7/17/13), at 1433; Arista User Manual, v. 4.11.1 (1/11/13), at 1179; Arista User Manual v. 4.10.3 (10/22/12), at 989; Arista User Manual v. 4.9.3.2 (5/3/12), at 748; ; Arista User Manual v. 4.8.2 (11/18/11), at 570.	Dkt. 419-10 at PDF p. 224

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Cisco IOS Interface and Hardware Component Command Reference (2011), at ISO-137.	**Display Values** • Inst. ID IS-IS Instance name. • System ID Identification value of the system listed in the Level 2 forwarding table. • Type Level 2 information. • Interface Interface through which the neighbor is reachable. • SNPA Subnetwork point of attachment (MAC address of the next hop). • State State of the adjacency: Up, Down, or INIT • Hold time Remaining hold time of the adjacency. • Area Address The address of the area. Arista User Manual v. 4.14.3F (Rev. 2) (10/2/2014), at 1702. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1459.	Dkt. 419-10 at PDF p. 225
Building the Address Table and Address Table Changes The device dynamically builds the address table by using the MAC source address of the frames received. When the device receives a frame for a MAC destination address not listed in its address table, it floods the frame to all LAN ports of the same VLAN except the port that received the frame. When the destination station replies, the device adds its relevant MAC source address and port ID to the address table. The device then forwards subsequent frames to a single LAN port without flooding all LAN ports. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 10.	**14.3 MAC Address Table** The switch maintains an MAC address table for switching frames efficiently between ports. The MAC address table contains static and dynamic MAC addresses. • Static MAC addresses are entered into the table through a CLI command. • Dynamic MAC addresses are entered into the table when the switch receives a frame whose source address is not listed in the MAC address table. The switch builds the table dynamically by referencing the source address of frames it receives. When the switch receives a frame, it associates the MAC address of the transmitting interface with the recipient VLAN. When a VLAN receives a frame for a MAC destination address not listed in the address table, the switch bridges the frame to all of the VLAN's ports except the recipient port. When the destination interface replies, the switch adds its MAC address to the MAC address table. The switch forwards subsequent frames with the destination address to the specified port. A multicast address can be associated with multiple ports. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 624. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 494; Arista User Manual, v. 4.11.1 (1/11/13), at 396-97; Arista User Manual v. 4.10.3 (10/22/12), at 328; Arista User Manual v. 4.9.3.2 (5/3/12), at 306.	Dkt. 419-10 at PDF p. 225

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
• Community VLAN—A community VLAN is a secondary VLAN that carries upstream traffic from the community ports to the promiscuous port gateways and to other host ports in the same community. You can configure multiple community VLANs in a private VLAN domain. The ports within one community can communicate, but these ports cannot communicate with ports in any other community or isolated VLAN in the private VLAN. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 54.	— Community Community VLAN ports carry traffic from host ports to the primary VLAN ports and to other host ports in the same community VLAN. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 763. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 611; Arista User Manual, v. 4.11.1 (1/11/13), at 467; Arista User Manual v. 4.10.3 (10/22/12), at 387; Arista User Manual v. 4.9.3.2 (5/3/12), at 307.	Dkt. 419-10 at PDF p. 226
• Protocol migration—For backward compatibility with 802.1D devices, 802.1w selectively sends 802.1D configuration BPDUs and TCN BPDUs on a per-port basis. When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which 802.1w BPDUs are sent), and 802.1w BPDUs are sent. While this timer is active, the device processes all BPDUs received on that port and ignores the protocol type. If the device receives an 802.1D BPDU after the port migration-delay timer has expired, it assumes that it is connected to an 802.1D device and starts using only 802.1D BPDUs. However, if the 802.1w device is using 802.1D BPDUs on a port and receives an 802.1w BPDU after the timer has expired, it restarts the timer and starts using 802.1w BPDUs on that port. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 100	The clear spanning-tree detected-protocols command forces MST ports to renegotiate with their neighbors. RSTP provides backward compatibility with 802.1D bridges as follows: • RSTP selectively sends 802.1D-configured BPDUs and Topology Change Notification (TCN) BPDUs on a per-port basis. • When a port initializes, the migration delay timer starts and RSTP BPDUs are transmitted. While the migration delay timer is active, the bridge processes all BPDUs received on that port. • If the bridge receives an 802.1D BPDU after a port's migration delay timer expires, the bridge assumes it is connected to an 802.1D bridge and starts using only 802.1D BPDUs. • When RSTP uses 802.1D BPDUs on a port and receives an RSTP BPDU after the migration delay expires, RSTP restarts the migration delay timer and resumes using RSTP BPDUs on that port. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 953. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 831; Arista User Manual, v. 4.11.1 (1/11/13), at 649; Arista User Manual v. 4.10.3 (10/22/12), at 563; Arista User Manual v. 4.9.3.2 (5/3/12), at 483; Arista User Manual v. 4.8.2 (11/18/11), at 357; Arista User Manual v. 4.7.3 (7/18/11), at 231.	Dkt. 419-10 at PDF p. 226

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Loop Guard Loop Guard helps prevent bridging loops that could occur because of a unidirectional link failure on a point-to-point link. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 176.	• Loop Guard: Prevents loops resulting from a unidirectional link failure on a point-to-point link. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 963. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 842; Arista User Manual, v. 4.11.1 (1/11/13), at 660; Arista User Manual v. 4.10.3 (10/22/12), at 574; Arista User Manual v. 4.9.3.2 (5/3/12), at 494; Arista User Manual v. 4.8.2 (11/18/11), at 368; Arista User Manual v. 4.7.3 (7/18/11), at 242.	Dkt. 419-10 at PDF p. 227
Rapid PVST+ achieves rapid transition to the forwarding state only on edge ports and point-to-point links. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 90.	RSTP only achieves rapid transition to forwarding state on edge ports and point-to-point links. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 964. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 843; Arista User Manual, v. 4.11.1 (1/11/13), at 661; Arista User Manual v. 4.10.3 (10/22/12), at 575; Arista User Manual v. 4.9.3.2 (5/3/12), at 494; Arista User Manual v. 4.8.2 (11/18/11), at 368; Arista User Manual v. 4.7.3 (7/18/11), at 242.	Dkt. 419-10 at PDF p. 227
Enabling Loop Guard on a root device has no effect but provides protection when a root device becomes a nonroot device. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 176.	Enabling loop guard on a root switch has no effect until the switch becomes a nonroot switch. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 966. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 844; Arista User Manual, v. 4.11.1 (1/11/13), at 662; Arista User Manual v. 4.10.3 (10/22/12), at 576; Arista User Manual v. 4.9.3.2 (5/3/12), at 496; Arista User Manual v. 4.8.2 (11/18/11), at 370; Arista User Manual v. 4.7.3 (7/18/11), at 244.	Dkt. 419-10 at PDF p. 227

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
• Enabling Loop Guard globally works only on point-to-point links. • Enabling Loop Guard per interface works on both shared and point-to-point links. • Root Guard forces a port to always be a designated port; it does not allow a port to become a root port. Loop Guard is effective only if the port is a root port or an alternate port. You cannot enable Loop Guard and Root Guard on a port at the same time. • Loop Guard has no effect on a disabled spanning tree instance or a VLAN. • Spanning tree always chooses the first operational port in the channel to send the BPDUs. If that link becomes unidirectional, Loop Guard blocks the channel, even if other links in the channel are functioning properly. • If you group a set of ports that are already blocked by Loop Guard to form a channel, spanning tree loses all the state information for those ports and the new channel port may obtain the forwarding state with a designated role. • If a channel is blocked by Loop Guard and the channel members go back to an individual link status, spanning tree loses all the state information. The individual physical ports may obtain the forwarding state with the designated role, even if one or more of the links that formed the channel are unidirectional. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 179.	Loop guard, when enabled globally, applies to all point-to-point ports. Loop guard is configurable on individual ports and applies to all STP instances of an enabled port. Loop-inconsistent ports transition to listening state when loop guard is disabled. Enabling loop guard on a root switch has no effect until the switch becomes a nonroot switch. When using loop guard: • Do not enable loop guard on portfast-enabled ports. • Loop guard is not functional on ports not connected to point-to-point links. • Loop guard has no effect on disabled spanning tree instances. Loop guard aspects on port channels include: • BPDUs are sent over the channel's first operational port. Loop guard blocks the channel if that link becomes unidirectional even when other channel links function properly. • Creating a new channel destroys state information for its component ports; new channels with loop-guard-enabled ports can enter forwarding state as a DP. • Dissembling a channel destroys its state information; component ports from a blocked channel can enter the forwarding state as DPs, even if the channel contained unidirectional links. • A unidirectional link on any port of a loop-guard-enabled channel blocks the entire channel until the affected port is removed or the link resumes bidirectional operation. Loop guard configuration commands include: • spanning-tree loopguard default command enables loop guard as a default on all switch ports. • spanning-tree guard control the loop guard setting on the configuration mode interface. This command overrides the default command for the specified interface. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 966. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 844; Arista User Manual, v. 4.11.1 (1/11/13), at 662; Arista User Manual v. 4.10.3 (10/22/12), at 576; Arista User Manual v. 4.9.3.2 (5/3/12), at 496; Arista User Manual v. 4.8.2 (11/18/11), at 370; Arista User Manual v. 4.7.3 (7/18/11), at 245.	Dkt. 419-10 at PDF p. 228

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
BPDU Guard Enabling BPDU Guard shuts down that interface if a BPDU is received. You can configure BPDU Guard at the interface level. When configured at the interface level, BPDU Guard shuts the port down as soon as the port receives a BPDU, regardless of the port type configuration. When you configure BPDU Guard globally, it is effective only on operational spanning tree edge ports. In a valid configuration, Layer 2 LAN edge interfaces do not receive BPDUs. A BPDU that is received by an edge Layer 2 LAN interface signals an invalid configuration, such as the connection of an unauthorized device. BPDU Guard, when enabled globally, shuts down all spanning tree edge ports when they receive a BPDU. BPDU Guard provides a secure response to invalid configurations, because you must manually put the Layer 2 LAN interface back in service after an invalid configuration. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 174-75.	20.3.4.3 BPDU Guard PortFast interfaces do not receive BPDUs in a valid configuration. BPDU Guard provides a secure response to invalid configurations by disabling ports when they receive a BPDU. Disabled ports differ from blocked ports in that they are re-enabled only through manual intervention. • When configured globally, BPDU Guard is enabled on ports in the operational portfast state. • When configured on an individual interface, BPDU Guard disables the port when it receives a BPDU, regardless of the port's portfast state. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 968. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 846; Arista User Manual, v. 4.11.1 (1/11/13), at 664-65; Arista User Manual v. 4.10.3 (10/22/12), at 578; Arista User Manual v. 4.9.3.2 (5/3/12), at 498; Arista User Manual v. 4.8.2 (11/18/11), at 372; Arista User Manual v. 4.7.3 (7/18/11), at 246.	Dkt. 419-10 at PDF p. 229
BPDU Filtering You can use BPDU Filtering to prevent the device from sending or even receiving BPDUs on specified ports. When configured globally, BPDU Filtering applies to all operational spanning tree edge ports. You should connect edge ports only to hosts, which typically drop BPDUs. If an operational spanning tree edge port receives a BPDU, it immediately returns to a normal spanning tree port type and moves through the regular transitions. In that case, BPDU Filtering is disabled on this port, and spanning tree resumes sending BPDUs on this port. In addition, you can configure BPDU Filtering by the individual interface. When you explicitly configure BPDU Filtering on a port, that port does not send any BPDUs and drops all BPDUs that it receives. You can effectively override the global BPDU Filtering setting on individual ports by configuring the specific interface. This BPDU Filtering command on the interface applies to the entire interface, whether the interface is trunking or not. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 175.	20.3.4.4 BPDU Filter BPDU filtering prevents the switch from sending or receiving BPDUs on specified ports. BPDU filtering is configurable on Ethernet and port channel interfaces. Ports with BPDU filtering enabled do not send BPDUs and drops inbound BPDUs. Enabling BPDU filtering on a port not connected to a host can result in loops as the port continues forwarding data while ignoring inbound BPDU packets. The spanning-tree bpdufilter command controls BPDU filtering on the configuration mode interface. BPDU filtering is disabled by default. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 968. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 846-47; Arista User Manual, v. 4.11.1 (1/11/13), at 665; Arista User Manual v. 4.10.3 (10/22/12), at 579; Arista User Manual v. 4.9.3.2 (5/3/12), at 498; Arista User Manual v. 4.8.2 (11/18/11), at 372; Arista User Manual v. 4.7.3 (7/18/11), at 246.	Dkt. 419-10 at PDF p. 229

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Bridge Assurance You can use Bridge Assurance to protect against certain problems that can cause bridging loops in the network. Specifically, you use Bridge Assurance to protect against a unidirectional link failure or other software failure and a device that continues to forward data traffic when it is no longer running the spanning tree algorithm. ✎ Note Bridge Assurance is supported only by Rapid PVST+ and MST. Bridge Assurance is enabled by default and can only be disabled globally. Also, Bridge Assurance can be enabled only on spanning tree network ports that are point-to-point links. Finally, both ends of the link must have Bridge Assurance enabled. If the device on one side of the link has Bridge Assurance enabled and the device on the other side either does not support Bridge Assurance or does not have this feature enabled, the connecting port is blocked. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 175.	**spanning-tree bridge assurance** The spanning-tree bridge assurance command enables bridge assurance on all ports with a port type of *network*. Bridge assurance protects against unidirectional link failure, other software failure, and devices that quit running a spanning tree algorithm. Bridge assurance is available only on spanning tree *network* ports on point-to-point links. Both ends of the link must have bridge assurance enabled. If the device on one side of the link has bridge assurance enabled and the device on the other side either does not support bridge assurance or does not have it enabled, the bridge assurance enabled port is blocked. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1002. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 880; Arista User Manual, v. 4.11.1 (1/11/13), at 698; Arista User Manual v. 4.10.3 (10/22/12), at 612; Arista User Manual v. 4.9.3.2 (5/3/12), at 531; Arista User Manual v. 4.8.2 (11/18/11), at 403; Arista User Manual v. 4.7.3 (7/18/11), at 252.	Dkt. 419-10 at PDF p. 230
• Root Guard—Root Guard prevents the port from becoming the root in an STP topology. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 6.	• Root guard prevents a port from becoming a root or blocked port. A root guard port that receives a superior BPDU transitions to the root-inconsistent (blocked) state. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1005. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 883; Arista User Manual, v. 4.11.1 (1/11/13), at 701; Arista User Manual v. 4.10.3 (10/22/12), at 615; Arista User Manual v. 4.9.3.2 (5/3/12), at 534; Arista User Manual v. 4.8.2 (11/18/11), at 406; Arista User Manual v. 4.7.3 (7/18/11), at 268.	Dkt. 419-10 at PDF p. 230

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
✎ Note Do not disable spanning tree on a VLAN unless all switches and bridges in the VLAN have spanning tree disabled. You cannot disable spanning tree on some switches and bridges in a VLAN and leave it enabled on other switches and bridges in the VLAN. This action can have unexpected results because switches and bridges with spanning tree enabled will have incomplete information regarding the physical topology of the network. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide (2011), at 108.	Important When disabling spanning tree on a VLAN, ensure that all switches and bridges in the network disable spanning tree for the same VLAN. Disabling spanning tree on a subset of switches and bridges in a VLAN may have unexpected results because switches and bridges running spanning tree will have incomplete information regarding the network's physical topology. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1023. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 901; Arista User Manual, v. 4.11.1 (1/11/13), at 719; Arista User Manual v. 4.10.3 (10/22/12), at 633; Arista User Manual v. 4.9.3.2 (5/3/12), at 550; Arista User Manual v. 4.8.2 (11/18/11), at 422; Arista User Manual v. 4.7.3 (7/18/11), at 264.	Dkt. 419-10 at PDF p. 231
The software elects a router as the IGMP querier on a subnet if it has the lowest IP address. As long as a router continues to receive query messages from a router with a lower IP address, it resets a timer that is based on its querier timeout value. If the querier timer of a router expires, it becomes the designated querier. If that router later receives a host query message from a router with a lower IP address, it drops its role as the designated querier and sets its querier timer again. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 20 .	The router with the lowest IP address on a subnet sends membership queries as the IGMP querier. When a router receives a membership query from a source with a lower IP address, it resets its query response timer. Upon timer expiry, the router begins sending membership queries If the router subsequently receives a membership query from a router with a lower IP address, it stops sending membership queries and resets the query response timer. Arista User Manual v. 4v. 4.14.3F – Rev. 2 (10/2/14), at 1779. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1505; Arista User Manual, v. 4.11.1 (1/11/13), at 1205; Arista User Manual v. 4.10.3 (10/22/12), at 999; Arista User Manual v. 4.9.3.2 (5/3/12), at 757; Arista User Manual v. 4.8.2 (11/18/11), at 579; Arista User Manual v. 4.7.3 (7/18/11), at 459; Arista User Manual v. 4.6.0 (12/22/2010), at 309	Dkt. 419-10 at PDF p. 231

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<table><tr><td>IGMP version</td><td>2</td></tr><tr><td>Startup query interval</td><td>30 seconds</td></tr><tr><td>Startup query count</td><td>2</td></tr><tr><td>Robustness value</td><td>2</td></tr><tr><td>Querier timeout</td><td>255 seconds</td></tr><tr><td>Query timeout</td><td>255 seconds</td></tr><tr><td>Query max response time</td><td>10 seconds</td></tr><tr><td>Query interval</td><td>125 seconds</td></tr><tr><td>Last member query response interval</td><td>1 second</td></tr><tr><td>Last member query count</td><td>2</td></tr><tr><td>Group membership timeout</td><td>260 seconds</td></tr><tr><td>Report link local multicast groups</td><td>Disabled</td></tr><tr><td>Enforce router alert</td><td>Disabled</td></tr><tr><td>Immediate leave</td><td>Disabled</td></tr></table>

Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 24. | ```
Current IGMP router version: 2
IGMP query interval: 125 seconds
IGMP max query response time: 100 deciseconds
Last member query response interval: 10 deciseconds
Last member query response count: 2
IGMP querier: 172.17.26.1
Robustness: 2
Require router alert: enabled
Startup query interval: 312 deciseconds
Startup query count: 2
General query timer expiry: 00:00:22
Multicast groups joined:
   239.255.255.250
```

Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1850.

See also Arista User Manual v. 4.12.3 (7/17/13), at 1558; Arista User Manual, v. 4.11.1 (1/11/13), at 1253; Arista User Manual v. 4.10.3 (10/22/12), at 1038; Arista User Manual v. 4.9.3.2 (5/3/12), at 796; Arista User Manual v. 4.8.2 (11/18/11), at 614; Arista User Manual v. 4.7.3 (7/18/11), at 491; Arista User Manual v. 4.6.0 (12/22/2010), at 337. | Dkt. 419-10 at PDF p. 232 |

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Anycast-RP Anycast-RP has two implementations: one uses Multicast Source Discovery Protocol (MSDP) and the other is based on *RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM)*. This section describes how to configure PIM Anycast-RP. You can use PIM Anycast-RP to assign a group of routers, called the Anycast-RP set, to a single RP address that is configured on multiple routers. The set of routers that you configure as Anycast-RPs is called the Anycast-RP set. This method is the only RP method that supports more than one RP per multicast group, which allows you to load balance across all RPs in the set. The Anycast RP supports all multicast groups. PIM register messages are sent to the closest RP and PIM join-prune messages are sent in the direction of the closest RP as determined by the unicast routing protocols. If one of the RPs goes down, unicast routing ensures these message will be sent in the direction of the next-closest RP. You must configue PIM on the loopback interface that is used for the PIM Anycast RP. For more information about PIM Anycast-RP, see *RFC 4610*. For information about configuring Anycast-RPs, see *Configuring a PIM Anycast-RP Set*. **PIM Register Messages** PIM register messages are unicast to the RP by designated routers (DRs) that are directly connected to multicast sources. The PIM register message has the following functions: • To notify the RP that a source is actively sending to a multicast group. • To deliver multicast packets sent by the source to the RP for delivery down the shared tree. The DR continues to send PIM register messages to the RP until it receives a Register-Stop message from the RP. The RP sends a Register-Stop message in either of the following cases: • The RP has no receivers for the multicast group being transmitted. • The RP has joined the SPT to the source but has not started receiving traffic from the source. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 68-69.	**Anycast-RP** PIM Anycast-RP defines a single RP address that is configured on multiple routers. An anycast-RP set consists of the routers configured with the same anycast-RP address. Anycast-RP provides redundancy protection and load balancing. The anycast-RP set supports all multicast groups. PIM register messages are unicast to the RP by designated routers (DRs) that are directly connected to multicast sources. The switch sends these messages and join-prune messages to the anycast-RP set member specified in the anycast-RP command. In a typical configuration, one command is required for each member of the anycast-RP set. The PIM register message has the following functions: • Notify the RP that a source is actively sending to a multicast group. • Deliver multicast packets sent by the source to the RP for delivery down the shared tree. The DR continues sending PIM register messages to the RP until it receives a Register-Stop message from the RP. The RP sends a Register-Stop message in either of the following cases: • The RP has no receivers for the multicast group being transmitted. • The RP has joined the SPT to the source but has not started receiving traffic from the source. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1874. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1580; Arista User Manual, v. 4.11.1 (1/11/13), at 1274; Arista User Manual v. 4.10.3 (10/22/12), at 1005-06; Arista User Manual v. 4.9.3.2 (5/3/12), at 763-64; Arista User Manual v. 4.8.2 (11/18/11), at 639; Arista User Manual v. 4.7.3 (7/18/11), at 514.	Dkt. 419-10 at PDF p. 233
Note Use the show ip mroute command to display the statistics for multicast route and prefixes. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 118 .	**Multicast Display Commands** To display the information in the multicast routing table, use the show ip mroute command. To display the MFIB table information, use the show ip mfib command. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1758. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1486; Arista User Manual, v. 4.11.1 (1/11/13), at 1188; Arista User Manual v. 4.10.3 (10/22/12), at 1012; Arista User Manual v. 4.9.3.2 (5/3/12), at 770; Arista User Manual v. 4.8.2 (11/18/11), at 589; Arista User Manual v. 4.7.3 (7/18/11), at 469; Arista User Manual v. 4.6.0 (12/22/2010), at 319.	Dkt. 419-10 at PDF p. 233

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
show ip mroute — Displays the contents of the IP multicast routing table. Cisco IOS IP Multicast Command Reference (July 16, 2005), at 12.	**Multicast Display Commands** To display the information in the multicast routing table, use the show ip mroute command. To display the MFIB table information, use the show ip mfib command. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1758 *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1486; Arista User Manual, v. 4.11.1 (1/11/13), at 1188; Arista User Manual v. 4.10.3 (10/22/12), at 1012; Arista User Manual v. 4.9.3.2 (5/3/12), at 770; Arista User Manual v. 4.8.2 (11/18/11), at 589; Arista User Manual v. 4.7.3 (7/18/11), at 469; Arista User Manual v. 4.6.0 (12/22/2010), at 319	Dkt. 419-10 at PDF p. 234
Command or Action — **Purpose** Step 4 — **Option** / **Description** — These commands configure IGMP snooping parameters. ip igmp snooping switch(config-vlan-config)# ip igmp snooping — Enables IGMP snooping for the current VLAN. The default is enabled. ip igmp snooping explicit-tracking switch(config-vlan-config)# ip igmp snooping explicit-tracking — Tracks IGMPv3 membership reports from individual hosts for each port on a per-VLAN basis. The default is enabled on all VLANs. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 139	The ip igmp snooping command controls the global snooping setting. The ip igmp snooping vlan command enables snooping on individual VLANs if snooping is globally enabled. IGMP snooping is enabled on all VLANs by default. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1780 *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1506; Arista User Manual, v. 4.11.1 (1/11/13), at 1206; Arista User Manual v. 4.10.3 (10/22/12), at 998; Arista User Manual v. 4.9.3.2 (5/3/12), at 756; Arista User Manual v. 4.8.2 (11/18/11), at 581; Arista User Manual v. 4.7.3 (7/18/11), at 461.	Dkt. 419-10 at PDF p. 234

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record

ip igmp snooping mrouter interface *interface*
`switch(config-vlan-config)#`
`ip igmp snooping mrouter`
`interface ethernet 2/1`
 	
Configures a static connection to a multicast router. The interface to the router must be in the selected VLAN. You can specify the interface by the type and the number, such as **ethernet** *slot/port*.
 	
Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 140.
	Specifying a Static Multicast Router Connection
	The **ip igmp snooping vlan mrouter** command statically configures a port that connects to a multicast router to join all multicast groups. The port to the router must be in the specified VLAN range.
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1780
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 1506; Arista User Manual, v. 4.11.1 (1/11/13), at 1206; Arista User Manual v. 4.10.3 (10/22/12), at 1003; Arista User Manual v. 4.9.3.2 (5/3/12), at 761; Arista User Manual v. 4.8.2 (11/18/11), at 584; Arista User Manual v. 4.7.3 (7/18/11), at 503; Arista User Manual v. 4.6.0 (12/22/2010), at 349.
		Dkt. 419-10 at PDF p. 234

Displaying IGMP Snooping Statistics Use the show ip igmp snooping statistics vlan command to display IGMP snooping statistics. You can see the virtual port channel (vPC) statistics in this output. Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 144	**show ip igmp statistics** The show ip igmp statistics command displays IGMP transmission statistics for the specified interface. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1867.	Dkt. 419-10 at PDF p. 235
SA Messages and Caching MSDP peers exchange Source-Active (SA) messages to propagate information about active sources. SA messages contain the following information: • Source address of the data source • Group address that the data source uses • IP address of the RP or the configured originator ID Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 148-49	35.2.2.1 Source Active Messages A Source Active (SA) message is a message that an RP creates and sends to MSDP peers when it learns of a new multicast source through a PIM register message. RPs that intend to originate or receive SA messages must establish MSDP peering with other RPs, either directly or through intermediate MSDP peers. An RP that is not a DR on a shared network should only originate SAs in response to register messages it receives from the DR. It does not originate SA's for directly connected sources in its domain. SA messages contain the following fields: • Source address of the data source. • Group address that receives data sent by the source. • IP address of the RP. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1912. Arista User Manual v. 4.12.3 (7/17/13), at 1618; Arista User Manual, v. 4.11.1 (1/11/13), at 1310.	Dkt. 419-10 at PDF p. 235

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
RFC 5059 — *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)* Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide (2012), at 174.	**34.3 Configuring PIM** The following sections describe the configuration of static RPs, dynamic RPs, and anycast-RPs. RP implementation is defined through the following RFCs: • RFC 5059: Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM). • RFC 6226: PIM Group-to-Rendezvous-Point Mapping. This section describes the following configuration tasks: • Section 34.3.1: Enabling PIM • Section 34.3.2: Rendezvous Points (RPs) • Section 34.3.3: Hello Messages • Section 34.3.4: Designated Router Election • Section 34.3.5: Join-Prune Messages Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1872. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 1578; Arista User Manual, v. 4.11.1 (1/11/13), at 1272; Arista User Manual v. 4.10.3 (10/22/12), at 1004; Arista User Manual v. 4.9.3.2 (5/3/12), at 762.	Dkt. 419-10 at PDF p. 236
Audience This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches. Cisco DCNM Fundamentals Guide, Release 6.x (2011), at lxi.	**Audience** This guide is for experienced network administrators who are responsible for configuring and maintaining Arista switches. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 41. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 35; Arista User Manual, v. 4.11.1 (1/11/13), at 29; Arista User Manual v. 4.10.3 (10/22/12), at 27; Arista User Manual v. 4.9.3.2 (5/3/12), at 23; Arista User Manual v. 4.8.2 (11/18/11), at 19; Arista User Manual v. 4.7.3 (7/18/11), at 17; Arista User Manual v. 4.6.0 (12/22/2010), at 13	Dkt. 419-10 at PDF p. 236

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Table 5-1 Channel Modes for Individual Links in a Port Channel **Channel Mode** / **Description** passive — LACP mode that places a port into a passive negotiating state in which the port responds to LACP packets that it receives but does not initiate LACP negotiation. active — LACP mode that places a port into an active negotiating state in which the port initiates negotiations with other ports by sending LACP packets. on — All static port channels (that are not running LACP) remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device displays an error message. You enable LACP on each channel by configuring the interface in that channel for the channel mode as either **active** or **passive**. When an LACP attempts to negotiate with an interface in the **on** state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group. The default port-channel mode is **on**. Interfaces Configuration Guide, Cisco DCNM for LAN, Release 6.x (2012), at 5-10	Parameters • *number* specifies a channel group ID. Values range from 1 through 1000. • LACP_MODE specifies the interface LACP mode. Values include: — mode on Configures interface as a static port channel, disabling LACP. The switch does not verify or negotiate port channel membership with other switches. — mode active Enables LACP on the interface in active negotiating state. The port initiates negotiations with other ports by sending LACP packets. — mode passive Enables LACP on the interface in a passive negotiating state. The port responds to LACP packets but cannot start LACP negotiations. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 469. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 403; Arista User Manual, v. 4.11.1 (1/11/13), at 336; Arista User Manual v. 4.10.3 (10/22/12), at 294; Arista User Manual v. 4.9.3.2 (5/3/12), at 278; Arista User Manual v. 4.8.2 (11/18/11), at 210; Arista User Manual v. 4.7.3 (7/18/11), at 424; Arista User Manual v. 4.6.0 (12/22/2010), at 271	Dkt. 419-10 at PDF p. 237
Table 6-1 Channel Modes for Individual Links in a Port Channel **Channel Mode** / **Description** passive — LACP mode that places a port into a passive negotiating state in which the port responds to LACP packets that it receives but does not initiate LACP negotiation. active — LACP mode that places a port into an active negotiating state in which the port initiates negotiations with other ports by sending LACP packets. on — All static port channels (that are not running LACP) remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device displays an error message. You enable LACP on each channel by configuring the interface in that channel for the channel mode as either **active** or **passive**. When an LACP attempts to negotiate with an interface in the **on** state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group. The default port-channel mode is **on**. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 6-10	• LACP_MODE specifies the interface LACP mode. Values include: — mode on Configures interface as a static port channel, disabling LACP. The switch does not verify or negotiate port channel membership with other switches. — mode active Enables LACP on the interface in active negotiating state. The port initiates negotiations with other ports by sending LACP packets. — mode passive Enables LACP on the interface in a passive negotiating state. The port responds to LACP packets but cannot start LACP negotiations. Arista User Manual v. 4.14.3F (Rev. 2) (October 2, 2014), at 469. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 403; Arista User Manual, v. 4.11.1 (1/11/13), at 336; Arista User Manual v. 4.10.3 (10/22/12), at 294; Arista User Manual v. 4.9.3.2 (5/3/12), at 278; Arista User Manual v. 4.8.2 (11/18/11), at 210; Arista User Manual v. 4.7.3 (7/18/11), at 424; Arista User Manual v. 4.6.0 (12/22/2010), at 271	Dkt. 419-10 at PDF p. 237

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Table 5-1 Channel Modes for Individual Links in a Port Channel **Channel Mode** / **Description** passive — LACP mode that places a port into a passive negotiating state in which the port responds to LACP packets that it receives but does not initiate LACP negotiation. active — LACP mode that places a port into an active negotiating state in which the port initiates negotiations with other ports by sending LACP packets. on — All static port channels (that are not running LACP) remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device displays an error message. You enable LACP on each channel by configuring the interface in that channel for the channel mode as either **active** or **passive**. When an LACP attempts to negotiate with an interface in the **on** state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group. The default port-channel mode is **on**. Interfaces Configuration Guide, Cisco DCNM for LAN, Release 5.x (2010), at 6-9	Parameters • *number* specifies a channel group ID. Values range from 1 through 1000. • LACP_MODE specifies the interface LACP mode. Values include: — mode on Configures interface as a static port channel, disabling LACP. The switch does not verify or negotiate port channel membership with other switches. — mode active Enables LACP on the interface in active negotiating state. The port initiates negotiations with other ports by sending LACP packets. — mode passive Enables LACP on the interface in a passive negotiating state. The port responds to LACP packets but cannot start LACP negotiations. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 469. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 403; Arista User Manual, v. 4.11.1 (1/11/13), at 336; Arista User Manual v. 4.10.3 (10/22/12), at 294; Arista User Manual v. 4.9.3.2 (5/3/12), at 278; Arista User Manual v. 4.8.2 (11/18/11), at 210; Arista User Manual v. 4.7.3 (7/18/11), at 424; Arista User Manual v. 4.6.0 (12/22/2010), at 271	Dkt. 419-10 at PDF p. 238
Table 5-1 Channel Modes for Individual Links in a Port Channel **Channel Mode** / **Description** passive — LACP mode that places a port into a passive negotiating state in which the port responds to LACP packets that it receives but does not initiate LACP negotiation. active — LACP mode that places a port into an active negotiating state in which the port initiates negotiations with other ports by sending LACP packets. on — All static port channels (that are not running LACP) remain in this mode. If you attempt to change the channel mode to active or passive before enabling LACP, the device displays an error message. You enable LACP on each channel by configuring the interface in that channel for the channel mode as either **active** or **passive**. When an LACP attempts to negotiate with an interface in the **on** state, it does not receive any LACP packets and becomes an individual link with that interface; it does not join the LACP channel group. The default port-channel mode is **on**. Interfaces Configuration Guide, Cisco DCNM for LAN, Release 4.x (2008), at 5-9	Parameters • *number* specifies a channel group ID. Values range from 1 through 1000. • LACP_MODE specifies the interface LACP mode. Values include: — mode on Configures interface as a static port channel, disabling LACP. The switch does not verify or negotiate port channel membership with other switches. — mode active Enables LACP on the interface in active negotiating state. The port initiates negotiations with other ports by sending LACP packets. — mode passive Enables LACP on the interface in a passive negotiating state. The port responds to LACP packets but cannot start LACP negotiations. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 469. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 403; Arista User Manual, v. 4.11.1 (1/11/13), at 336; Arista User Manual v. 4.10.3 (10/22/12), at 294; Arista User Manual v. 4.9.3.2 (5/3/12), at 278; Arista User Manual v. 4.8.2 (11/18/11), at 210; Arista User Manual v. 4.7.3 (7/18/11), at 424; Arista User Manual v. 4.6.0 (12/22/2010), at 271	Dkt. 419-10 at PDF p. 238

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Note For information about configuring port channels and the Link Aggregation Control Protocol (LACP), see Chapter 5, "Configuring Port Channels." Interfaces Configuration Guide, Cisco DCNM for LAN, Release 6.x (2012), at 6-2	**Port Channels and LACP** This chapter describes channel groups, port channels, port channel interfaces, and the Link Aggregation Control Protocol (LACP). This chapter contains the following sections: Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 469. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 391; Arista User Manual, v. 4.11.1 (1/11/13), at 329; Arista User Manual v. 4.10.3 (10/22/12), at 287; Arista User Manual v. 4.9.3.2 (5/3/12), at 271; Arista User Manual v. 4.8.2 (11/18/11), at 203.	Dkt. 419-10 at PDF p. 239
Note For information about configuring port channels and the Link Aggregation Control Protocol (LACP), see Chapter 5, "Configuring Port Channels." Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 7-1	**Port Channels and LACP** This chapter describes channel groups, port channels, port channel interfaces, and the Link Aggregation Control Protocol (LACP). This chapter contains the following sections: Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 469. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 391; Arista User Manual, v. 4.11.1 (1/11/13), at 329; Arista User Manual v. 4.10.3 (10/22/12), at 287; Arista User Manual v. 4.9.3.2 (5/3/12), at 271; Arista User Manual v. 4.8.2 (11/18/11), at 203.	Dkt. 419-10 at PDF p. 239

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Note For information about configuring port channels and the Link Aggregation Control Protocol (LACP), see Chapter 5, "Configuring Port Channels." Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 5.x (2010), at 7-1	**Port Channels and LACP** This chapter describes channel groups, port channels, port channel interfaces, and the Link Aggregation Control Protocol (LACP). This chapter contains the following sections: Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 469. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 391; Arista User Manual, v. 4.11.1 (1/11/13), at 329; Arista User Manual v. 4.10.3 (10/22/12), at 287; Arista User Manual v. 4.9.3.2 (5/3/12), at 271; Arista User Manual v. 4.8.2 (11/18/11), at 203.	Dkt. 419-10 at PDF p. 240
Loopback Interfaces A loopback interface is a virtual interface with a single endpoint that is always up. Any packet transmitted over a loopback interface is immediately received by this interface. Loopback interfaces emulate a physical interface. You can configure up to 1024 loopback interfaces per VDC, numbered 0 to 1023. Interfaces Configuration Guide, Cisco DCNM for LAN, Release 6.x (2012), at 4-4 .	14.4.4 Loopback Ports A loopback interface is a virtual network interface implemented in software and does not connect to any hardware. Traffic sent to the loopback interface is immediately received on the sending interface. The switch provides loopback configuration mode for creating loopback interfaces and modifying their operating parameters. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 631. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 500; Arista User Manual, v. 4.11.1 (1/11/13), at 397; Arista User Manual v. 4.10.3 (10/22/12), at 329.	Dkt. 419-10 at PDF p. 240
Loopback Interfaces A loopback interface is a virtual interface with a single endpoint that is always up. Any packet transmitted over a loopback interface is immediately received by this interface. Loopback interfaces emulate a physical interface. You can configure up to 1024 loopback interfaces per VDC, numbered 0 to 1023. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x (2013), at 4-4	14.4.4 Loopback Ports A loopback interface is a virtual network interface implemented in software and does not connect to any hardware. Traffic sent to the loopback interface is immediately received on the sending interface. The switch provides loopback configuration mode for creating loopback interfaces and modifying their operating parameters. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 631. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 500; Arista User Manual, v. 4.11.1 (1/11/13), at 397; Arista User Manual v. 4.10.3 (10/22/12), at 329.	Dkt. 419-10 at PDF p. 240

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
Loopback Interfaces A loopback interface is a virtual interface with a single endpoint that is always up. Any packet transmitted over a loopback interface is immediately received by this interface. Loopback interfaces emulate a physical interface. You can configure up to 1024 loopback interfaces per VDC, numbered 0 to 1023. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 5.x (2010), at 4-4	14.4.4 Loopback Ports A loopback interface is a virtual network interface implemented in software and does not connect to any hardware. Traffic sent to the loopback interface is immediately received on the sending interface. The switch provides loopback configuration mode for creating loopback interfaces and modifying their operating parameters. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 631. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 500; Arista User Manual, v. 4.11.1 (1/11/13), at 397; Arista User Manual v. 4.10.3 (10/22/12), at 329.	Dkt. 419-10 at PDF p. 241
Loopback Interfaces A loopback interface is a virtual interface with a single endpoint that is always up. Any packet transmitted over a loopback interface is immediately received by this interface. Loopback interfaces emulate a physical interface. You can configure up to 1024 loopback interfaces per VDC, numbered 0 to 1023. Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 4.x (2010), at 4-3	14.4.4 Loopback Ports A loopback interface is a virtual network interface implemented in software and does not connect to any hardware. Traffic sent to the loopback interface is immediately received on the sending interface. The switch provides loopback configuration mode for creating loopback interfaces and modifying their operating parameters. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 631. *See also* Arista User Manual v. 4.12.3 (7/17/13), at 500; Arista User Manual, v. 4.11.1 (1/11/13), at 397; Arista User Manual v. 4.10.3 (10/22/12), at 329.	Dkt. 419-10 at PDF p. 241
Configuring a Maximum Number of MAC Addresses You can configure the maximum number of MAC addresses that can be learned or statically configured on interfaces that belong to a port profile. Interfaces Configuration Guide, Cisco DCNM for LAN, Release 6.x (2012), at 10-22	**Port Security Configuration** MAC security restricts input to a switched port by limiting the number and identity of MAC addresses that can access the port. MAC address security is enabled by switchport port-security. Ports with MAC security enabled restrict traffic to a limited number of hosts, as determined by their MAC addresses. The maximum number of MAC addresses that can be assigned to an interface is configured by switchport port-security maximum. The default MAC address limit on an interface where port security is enabled is one. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 632. *See also* Arista User Manual v. 4.13.6F (4/14/2014), at 624; Arista User Manual v. 4.12.3 (7/17/13), at 501; Arista User Manual, v. 4.11.1 (1/11/13), at 405; Arista User Manual v. 4.10.3 (10/22/12), at 336.	Dkt. 419-10 at PDF p. 241

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
By default, an interface can have only one secure MAC address. You can configure the maximum number of MAC addresses permitted per interface or per VLAN on an interface. Maximums apply to secure MAC addresses learned by any method: dynamic, sticky, or static. ICisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 6.x (2013), at 507	**Port Security Configuration** MAC security restricts input to a switched port by limiting the number and identity of MAC addresses that can access the port. MAC address security is enabled by switchport port-security. Ports with MAC security enabled restrict traffic to a limited number of hosts, as determined by their MAC addresses. The maximum number of MAC addresses that can be assigned to an interface is configured by switchport port-security maximum. The default MAC address limit on an interface where port security is enabled is one. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 632. *See also* Arista User Manual v. 4.13.6F (4/14/2014), at 624; Arista User Manual v. 4.12.3 (7/17/13), at 501; Arista User Manual, v. 4.11.1 (1/11/13), at 405; Arista User Manual v. 4.10.3 (10/22/12), at 336.	Dkt. 419-10 at PDF p. 242
By default, an interface can have only one secure MAC address. You can configure the maximum number of MAC addresses permitted per interface or per VLAN on an interface. Maximums apply to secure MAC addresses learned by any method: dynamic, sticky, or static. Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 5.x (2010), at 177	**Port Security Configuration** MAC security restricts input to a switched port by limiting the number and identity of MAC addresses that can access the port. MAC address security is enabled by switchport port-security. Ports with MAC security enabled restrict traffic to a limited number of hosts, as determined by their MAC addresses. The maximum number of MAC addresses that can be assigned to an interface is configured by switchport port-security maximum. The default MAC address limit on an interface where port security is enabled is one. Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 632. *See also* Arista User Manual v. 4.13.6F (4/14/2014), at 624; Arista User Manual v. 4.12.3 (7/17/13), at 501; Arista User Manual, v. 4.11.1 (1/11/13), at 405; Arista User Manual v. 4.10.3 (10/22/12), at 336.	Dkt. 419-10 at PDF p. 242

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record

This example shows how to return to EXEC mode from global configuration mode:
`switch(config)# end`
`switch#`
 	
This example shows how to return to EXEC mode from interface configuration mode:
`switch(config-if)# end`
`switch#`
 	
Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2013), at FND-44
	• To return to Privileged EXEC mode from any configuration mode, type end or Ctrl-Z.
	`switch(config-if-Et24)#<Ctrl-z>`
	`switch#`
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 120.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 99; Arista User Manual, v. 4.11.1 (1/11/13), at 69; Arista User Manual v. 4.10.3 (10/22/12), at 61; Arista User Manual v. 4.9.3.2 (5/3/12), at 57; Arista User Manual v. 4.8.2 (11/18/11), at 52; Arista User Manual v. 4.7.3 (7/18/11), at 47; Arista User Manual v. 4.6.0 (12/22/2010), at 41
		Dkt. 419-10 at PDF p. 243

This example shows how to return to EXEC mode from global configuration mode:
`switch(config)# end`
`switch#`
 	
This example shows how to return to EXEC mode from interface configuration mode:
`switch(config-if)# end`
`switch#`
 	
Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2010), at FND-37
	• To return to Privileged EXEC mode from any configuration mode, type end or Ctrl-Z.
	`switch(config-if-Et24)#<Ctrl-z>`
	`switch#`
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 120.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 99; Arista User Manual, v. 4.11.1 (1/11/13), at 69; Arista User Manual v. 4.10.3 (10/22/12), at 61; Arista User Manual v. 4.9.3.2 (5/3/12), at 57; Arista User Manual v. 4.8.2 (11/18/11), at 52; Arista User Manual v. 4.7.3 (7/18/11), at 47; Arista User Manual v. 4.6.0 (12/22/2010), at 41
		Dkt. 419-10 at PDF p. 243

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record

Note ✎ The reload command does not save the running configuration. Use the copy running-config startup-config command to save the current configuration on the device.
 	
Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2013), at FND-105
	Step 8 Type write memory (or copy running-config startup-config) to save the new configuration to the *startup-config* file. See Section 3.5.4: Saving the Running Configuration Settings.
	
	`switch# write memory`
	`switch#`
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 60.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 52; Arista User Manual, v. 4.11.1 (1/11/13), at 44; Arista User Manual v. 4.10.3 (10/22/12), at 38; Arista User Manual v. 4.9.3.2 (5/3/12), at 34; Arista User Manual v. 4.8.2 (11/18/11), at 30; Arista User Manual v. 4.7.3 (7/18/11), at 28; Arista User Manual v. 4.6.0 (12/22/2010), at 25
		Dkt. 419-10 at PDF p. 244

Note ✎ The reload command does not save the running configuration. Use the copy running-config startup-config command to save the current configuration on the device.
 	
Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2010), at FND-84
	Step 8 Type write memory (or copy running-config startup-config) to save the new configuration to the *startup-config* file. See Section 3.5.4: Saving the Running Configuration Settings.
	
	`switch# write memory`
	`switch#`
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 60.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 52; Arista User Manual, v. 4.11.1 (1/11/13), at 44; Arista User Manual v. 4.10.3 (10/22/12), at 38; Arista User Manual v. 4.9.3.2 (5/3/12), at 34; Arista User Manual v. 4.8.2 (11/18/11), at 30; Arista User Manual v. 4.7.3 (7/18/11), at 28; Arista User Manual v. 4.6.0 (12/22/2010), at 25
		Dkt. 419-10 at PDF p. 244

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record

This example shows how to display commands related to Open Shortest Path First (OSPF) available in the loopback interface command mode:
 	
`switch(config)# interface loopback 0`
`switch(config-if)# show cli list ospf`
`MODE if-loopback`
`no ip ospf network point-to-point`
`no ip ospf network`
 	
Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2013), at FND-126
	Command Syntax
	
	`ip ospf network point-to-point`
	`no ip ospf network`
	`default ip ospf network`
	
	Examples
	
	• These commands configure Ethernet interface 10 as a point-to-point link.
	
	`switch(config)#interface ethernet 10`
	`switch(config-if-Et10)#ip ospf network point-to-point`
	`switch(config-if-Et10)#`
	
	• This command restores Ethernet interface 10 as a broadcast link.
	
	`switch(config-if-Et10)#no ip ospf network`
	`switch(config-if-Et10)#`
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1432.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 1219; Arista User Manual, v. 4.11.1 (1/11/13), at 976; Arista User Manual v. 4.10.3 (10/22/12), at 806; Arista User Manual v. 4.9.3.2 (5/3/12), at 692; Arista User Manual v. 4.8.2 (11/18/11), at 465; Arista User Manual v. 4.7.3 (7/18/11), at 338.
		Dkt. 419-10 at PDF p. 245

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record

This example shows how to display commands related to Open Shortest Path First (OSPF) available in the loopback interface command mode:
 	
`switch(config)# interface loopback 0`
`switch(config-if)# show cli list ospf`
`MODE if-loopback`
`no ip ospf network point-to-point`
`no ip ospf network`
 	
Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2010), at FND-105
	Command Syntax
	
	`ip ospf network point-to-point`
	`no ip ospf network`
	`default ip ospf network`
	
	Examples
	
	• These commands configure Ethernet interface 10 as a point-to-point link.
	
	`switch(config)#interface ethernet 10`
	`switch(config-if-Et10)#ip ospf network point-to-point`
	`switch(config-if-Et10)#`
	
	• This command restores Ethernet interface 10 as a broadcast link.
	
	`switch(config-if-Et10)#no ip ospf network`
	`switch(config-if-Et10)#`
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 1432.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 1219; Arista User Manual, v. 4.11.1 (1/11/13), at 976; Arista User Manual v. 4.10.3 (10/22/12), at 806; Arista User Manual v. 4.9.3.2 (5/3/12), at 692; Arista User Manual v. 4.8.2 (11/18/11), at 465; Arista User Manual v. 4.7.3 (7/18/11), at 338.
		Dkt. 419-10 at PDF p. 246

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record

show startup-config
 	
To display the startup configuration use the show **startup-config** command.
 	
show startup-config [exclude *component-list*]
 	
Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2013), at FND-154.
	Example
	• Type show startup-config to display the startup configuration file. The response in the example is truncated to display only the ip route configured in Admin Username (page 58).
	
	`switch#show startup-config`
	`! Command: show startup-config`
	`! Startup-config last modified at  Wed Feb 19 08:34:31 2014 by admin`
	`!`
	`<-------OUTPUT OMITTED FROM EXAMPLE-------->`
	`!`
	`ip route 0.0.0.0/0 192.0.2.1`
	`!`
	`<-------OUTPUT OMITTED FROM EXAMPLE-------->`
	`end`
	`switch#`
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 123.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 102; Arista User Manual, v. 4.11.1 (1/11/13), at 72; Arista User Manual v. 4.10.3 (10/22/12), at 65; Arista User Manual v. 4.9.3.2 (5/3/12), at 59; Arista User Manual v. 4.8.2 (11/18/11), at 54; Arista User Manual v. 4.7.3 (7/18/11), at 49.
		Dkt. 419-10 at PDF p. 247

show startup-config
 	
To display the startup configuration use the show **startup-config** command.
 	
show startup-config [exclude *component-list*]
 	
Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference (2010), at FND-125.
	Example
	• Type show startup-config to display the startup configuration file. The response in the example is truncated to display only the ip route configured in Admin Username (page 58).
	
	`switch#show startup-config`
	`! Command: show startup-config`
	`! Startup-config last modified at  Wed Feb 19 08:34:31 2014 by admin`
	`!`
	`<-------OUTPUT OMITTED FROM EXAMPLE-------->`
	`!`
	`ip route 0.0.0.0/0 192.0.2.1`
	`!`
	`<-------OUTPUT OMITTED FROM EXAMPLE-------->`
	`end`
	`switch#`
	
	Arista User Manual v. 4.14.3F – Rev. 2 (10/2/14), at 123.
	
	See also Arista User Manual v. 4.12.3 (7/17/13), at 102; Arista User Manual, v. 4.11.1 (1/11/13), at 72; Arista User Manual v. 4.10.3 (10/22/12), at 65; Arista User Manual v. 4.9.3.2 (5/3/12), at 59; Arista User Manual v. 4.8.2 (11/18/11), at 54; Arista User Manual v. 4.7.3 (7/18/11), at 49.
		Dkt. 419-10 at PDF p. 248